

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BOARD OF PATENT APPEALS AND INTERFERENCES**

In Re Application of:

Wang *et al.*

Serial No.: 10/729,804

Filed: December 5, 2003

Confirmation No.: 2268

Group Art Unit: 2451

Examiner: Madamba, Glenford J.

TKHR Ref: 250338-1500

Client Ref: S-296

For: **CLASS-BASED RATE CONTROL USING MULTI-THRESHOLD LEAKY
BUCKET**

APPEAL BRIEF PURSUANT TO 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner of Patents and Trademarks
P.O. Box 1450
Alexandria, Virginia 22313-1450

Sir:

This Appeal Brief under 37 C.F.R. § 41.37 is submitted in response to the Notice of Panel Decision from Pre-Appeal Brief Review mailed November 24, 2009 and in further support of the Notice of Appeal filed on July 14, 2009, responding to the Office Action mailed April 14, 2009.

I. REAL PARTY IN INTEREST

The real party in interest is Conexant Systems Inc., having a place of business at 4000 MacArthur Blvd., Newport Beach, California 92660.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

III. STATUS OF THE CLAIMS

Claims 1-27 were rejected by the non-final Office Action mailed April 14, 2009 and are the subject of this appeal. No claims have been allowed.

IV. STATUS OF AMENDMENTS

Claims 1-27 were rejected by the non-final Office Action mailed April 14, 2009. Claims 1, 16, and 24 were amended during prosecution. No claims were amended after the non-final Office Action mailed April 14, 2009. Claims 1-27 are the subject of this appeal. The claims in the attached Claims Appendix reflect the present state of Appellants' claims.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The claimed inventions are summarized below with reference numerals and references to the written description ("specification") and drawings. The subject matter described in the following appears in the original disclosure at least where indicated, and may further appear in other places within the original disclosure.

An embodiment of the invention, such as the one defined by claim 1, is directed to an egress rate controller (see *e.g.*, FIG. 1, component 200, and page 8, line 30) monitoring content traffic transmitted from an edge network node of a packet-switched communications network node (see *e.g.*, page 9, lines 24-28). In accordance with such embodiments, the egress rate controller comprises a leaky bucket (see *e.g.*, page 9, lines 3-5) having an initial maximum number of tokens (see *e.g.*, page 9, lines 3-7) which decreases as packets are received in an associated output buffer at a reception token rate for transmission (see *e.g.*, page 9, lines 14-19), wherein a size of the leaky bucket is less

than or equal to a size of the associated output buffer (see *e.g.*, page 9, lines 8-10). The egress rate controller further comprises a plurality of token availability threshold level registers (see *e.g.*, page 9, line 1) specifying a corresponding plurality of token amounts defining token availability regions (see *e.g.*, page 10, lines 13-18), and a packet transmission suppression controller (see *e.g.*, FIGS. 1 and 2, component 204, and page 10, line 9) selectively suppressing transmission of a packet (see *e.g.*, page 9, lines 8-13) having a traffic class association based on a current token availability level being within a token availability region specifying transmission suppression of packets of the traffic class (see *e.g.*, page 10, lines 23-28).

An embodiment of the invention, such as the one defined by claim 9, is directed to an ingress rate controller (see *e.g.*, FIG. 1, component 300, and page 11, line 25) monitoring content traffic received at an edge network node of a packet-switched communications network node. In accordance with such embodiments, the ingress rate controller comprises a leaky bucket (see *e.g.*, page 12, lines 3-10) having an initial maximum number of tokens which decreases as packets received at a reception token rate are accepted (see *e.g.*, page 12, lines 15-19), a plurality of token availability threshold level registers specifying a corresponding plurality of token amounts defining token availability regions (see *e.g.*, page 12, lines 15-22), a plurality of packet discard probability registers (see *e.g.*, page 13, line 20), each packet discard probability register specifying a probability with which packets of a specific traffic class are to be dropped when a current token availability level is within a token availability region (see *e.g.*, page 13, lines 20-22), and a packet acceptance controller (see *e.g.*, page 13, lines 26-29) selectively randomly discarding packets having a traffic class association based on the current token availability

level being within a token availability region specifying random packet discard of packets of the traffic class (see *e.g.*, page 9, line 26 to page 14, line 3).

An embodiment of the invention, such as the one defined by claim 16, is directed to a method of effecting egress rate control (see *e.g.*, page 8, lines 3-4). In accordance with such embodiments, the method comprises selectively suppressing packet transmission (see *e.g.*, page 9, lines 3-4) for a packet of a particular traffic class when a current token availability level of a leaky bucket tracking packet transmissions is between two token availability threshold levels of a plurality of token availability threshold levels (see *e.g.*, page 10, lines 15-16), wherein the token availability threshold levels correspond to predetermined egress rate control responses (see *e.g.*, page 10, lines 15-23) to bandwidth utilization with respect to packet traffic classes (see *e.g.*, page 10, lines 25-28).

An embodiment of the invention, such as the one defined by claim 24, is directed to a method of effecting ingress rate control (see *e.g.*, page 12, lines 3-5). In accordance with such embodiments, the method comprises selectively randomly discarding packets of a particular traffic class when a current token availability level of a leaky bucket tracking packets is between two token availability threshold levels of a plurality of token availability threshold levels (see *e.g.*, page 13, lines 10-15), wherein the token availability threshold levels correspond to predetermined ingress rate control responses (see *e.g.*, page 13, lines 13-15) to bandwidth utilization with respect to packet traffic classes (see *e.g.*, page 13, lines 14-15).

VI. GROUND S OF REJECTION TO BE REVIEWED ON APPEAL

The following grounds of rejection are to be reviewed on appeal.

A. Claims 1-8, 11 and 17-23 are rejected under 35 U.S.C. 103(a) as allegedly being

unpatentable over *Carter, et al.* (U.S. Pub. No. 2003/0035374, hereinafter, “*Carter*”) in view of *Patel, et al.* (U.S. Patent No. 7,126,913, hereinafter “*Patel*”), and further in view of *Elwalid, et al.* (U.S. Patent No. 5,978,356, hereinafter “*Elwalid*”).

- B. Claim 16 is rejected under 35 U.S.C. 103(a) as allegedly being unpatentable over *Carter* in view of *Patel* further in view of *Lee, et al.* (U.S. Patent No. 7,349,403, hereinafter “*Lee*”).
- C. Claims 9-10, 12-15 and 24-27 are rejected under 35 U.S.C. 103(a) as allegedly being unpatentable over *Carter, Patel* and *Gracon et al.* (U.S. Patent No. 6,987,732, hereinafter “*Gracon*”), further in view of *Lee*.

VII. ARGUMENT

A. Rejection of Claims 1-8, 11, and 17-23 under 35 U.S.C. §103(a): *Carter, Patel, and Elwalid*

1. Independent Claim 1

The non-final Office Action mailed April 14, 2009 rejected claims 1-8 under 35 U.S.C. §103(a) as allegedly being unpatentable over *Carter* in view of *Patel* in further view of *Elwalid*. For at least the reasons set forth below, Appellants submit that these rejections should be overturned.

Appellants’ claim 1 provides as follows:

1. An egress rate controller monitoring content traffic transmitted from an edge network node of a packet-switched communications network node comprising:
 - a. a leaky bucket having an initial maximum number of tokens which decreases as packets are received in an associated output buffer at a reception token rate for transmission, **wherein a size of the leaky bucket is less than or equal to a size of the associated output buffer;**
 - b. a plurality of token availability threshold level registers specifying a corresponding plurality of token amounts defining

token availability regions; and

c. a packet transmission suppression controller selectively suppressing transmission of a packet having a traffic class association based on a current token availability level being within a token availability region specifying transmission suppression of packets of the traffic class.

(Emphasis added). In rejecting claim 1, the Office Action now cites the *Elwalid* reference in place of the previously cited *Aukia* reference (U.S. Patent No. 6,594,268) and alleges that *Elwalid* discloses the features “wherein a size of the leaky bucket is less than or equal to a size of the associated output buffer” and “based on a current token availability level being within a token availability region specifying transmission suppression of packets of the traffic class.” (Office Action mailed April 14, 2009, pages 5-6). In alleging that *Elwalid* discloses the feature “wherein a size of the leaky bucket is less than or equal to a size of the associated output buffer,” the Office Action refers to the parameters related to the Dual Leaky Bucket 16 taught by *Elwalid*. Specifically, the Office Action refers to the three parameters r , B_T , and P . Furthermore, the Examiner emphasizes the fact that the parameter B_T relates to the number of tokens that the token buffer 30 can hold (col. 5, line 16: “The token buffer 30 is capable of holding B_T tokens.”) and apparently equates this to the “size of the leaky bucket” in claim 1 as the Office Action alleges “*wherein ‘ B_T ’ is the leaky bucket ‘Token Buffer Size’.*” (Office Action mailed April 14, 2009, page 5). It is unclear, however, how the cited passage in *Elwalid* teaches the feature “wherein a size of the leaky bucket is less than or equal to a size of the associated output buffer.”

While the cited passage in *Elwalid* describes a rate r (the rate at which tokens are supplied to the token buffer 30) and that P (the peak rate) is greater than r , nowhere does *Elwalid* teach that B_T (allegedly the “size of the leaky bucket”) is less than or equal

to a size of the associated output buffer, as recited in claim 1. Appellants note that the parameters P and r relate to token rates. If anything, the P line token buffer 38 taught by *Elwalid* appears to be more relevant to the output buffer size limitation in claim 1. However, it is unclear what the relationship is between B_T and the size of the P line token buffer 38. In this regard, the newly cited *Elwalid* reference fails to disclose this feature in claim 1. For at least this reason, Appellants respectfully submit that independent claim 1 patently defines over *Carter* in view of *Patel* further in view of *Elwalid*.

2. Dependent Claims 2-8

Appellants submit that dependent claims 2-8 are allowable for at least the reason that these claims depend from an allowable independent claim. *See, e.g., In re Fine*, 837 F. 2d 1071 (Fed. Cir. 1988).

3. Dependent Claim 11

Claim 11 is rejected under 35 U.S.C. 103(a) as allegedly being unpatentable over *Carter* in view of *Patel* in further view of *Elwalid*. As set forth above, Appellants submit that independent claim 9, from which claim 11 depends, is patentable over *Carter*, *Patel*, *Gracon*, and *Lee*. Furthermore, the *Elwalid* reference fails to address the deficiencies expressed above for these references. As such, Appellants submit that independent claim 9 is patentable over the combination of *Carter* in view of *Patel* in further view of *Elwalid*. Accordingly, dependent claim 11 is allowable for at least the reason that this claim depends from an allowable independent claim. *See, e.g., In re Fine*, 837 F. 2d 1071 (Fed. Cir. 1988).

4. Dependent Claims 17-23

Claims 17-23 are rejected under 35 U.S.C. 103(a) as allegedly being unpatentable over *Carter* in view of *Patel* in further view of *Elwalid*. As set forth below, Appellants submit that independent claim 16, from which claims 17-23 depend, are patentable over *Carter*, *Patel*, and *Lee*. Furthermore, *Elwalid* fails to address the deficiencies expressed above for these references. As such, Appellants submit that independent claim 16 is patentable over the combination of *Carter* in view of *Patel* in further view of *Elwalid*. Accordingly, dependent claims 17-23 are allowable for at least the reason that these claims depend from an allowable independent claim. See, e.g., *In re Fine*, 837 F. 2d 1071 (Fed. Cir. 1988).

B. Rejection of Claim 16 under 35 U.S.C. §103(a): *Carter*, *Patel*, and *Lee*

The non-final Office Action mailed April 14, 2009 rejected claim 16 under 35 U.S.C. §103(a) as allegedly being unpatentable over *Carter* in view of *Patel* in further view of *Lee*. For at least the reasons set forth below, Appellants submit that these rejections should be overturned.

Appellants' claim 16 provides as follows:

16. A method of effecting egress rate control comprising the step of:
selectively suppressing packet transmission for a packet of a particular traffic class when a current token availability level of a leaky bucket tracking packet transmissions is between two token availability threshold levels of a plurality of token availability threshold levels, **wherein the token availability threshold levels correspond to predetermined egress rate control responses** to bandwidth utilization with respect to packet traffic classes.

(Emphasis added). Claim 16 recites "wherein the token availability threshold levels correspond to predetermined egress rate control responses to bandwidth utilization with respect to packet traffic classes." The Office Action equates these features with

“determining whether ‘average usage of a class to which a flow belongs’ is equal to, less than, or greater than a minimum/maximum threshold. [Figs. 1-2]” as taught by Lee. (Office Action mailed April 14, 2009, page 11). The Office Action further refers to the accepting or discarding of incoming information elements depicted in FIGS. 34-41. Claim 16 explicitly recites “token availability threshold levels” in addition to “predetermined egress rate control responses.” Appellants respectfully submit that these elements are not taught by Lee.

Claim 15 describes token availability threshold levels in the context of leaky bucket tracking packet transmissions (“when a current token availability level of a leaky bucket tracking packets is between two token availability threshold levels of a plurality of token availability threshold levels”). Even assuming, for the sake of argument, that the minimum and maximum thresholds cited in the Office Action correspond with the token availability threshold levels in claim 16, for example, Lee fails to disclose or suggest token availability threshold levels corresponding to predetermined egress rate control responses to bandwidth utilization. In FIG. 34, the compare unit 966 compares the average “information segment storage unit” occupancy of a particular class using the average occupancy counter for that class with the maximum number of occupied “information segment storage unit” rows and the minimum number of occupied “information segment storage unit” rows for that class. Based on this, the multiplexer 964 selects as its output a particular one of the inputs (e.g., “always discard”). Appellants submit that is not equivalent to the predetermined egress rate control responses defined in claim 16.

C. Rejection of Claims 9-10, 12-15, and 24-27 under 35 U.S.C. §103(a): Carter, Patel, Gracon, and Lee

1. Independent Claim 9

The non-final Office Action mailed April 14, 2009 rejected claims 9-10 and 12-15 under 35 U.S.C. §103(a) as allegedly being unpatentable over *Carter* in view of *Patel*, and *Gracon* in further view of *Lee*. For at least the reasons set forth below, Appellants submit that these rejections should be overturned.

Appellants' claim 9 provides as follows:

9. An ingress rate controller monitoring content traffic received at an edge network node of a packet-switched communications network node comprising:
- a. a leaky bucket having an initial maximum number of tokens which decreases as packets received at a reception token rate are accepted;
 - b. a plurality of token availability threshold level registers specifying a corresponding plurality of token amounts defining token availability regions;
 - c. a plurality of packet discard probability registers, each packet discard probability register specifying a probability with which packets of a specific traffic class are to be dropped when a current token availability level is within a token availability region,** and
 - d. a packet acceptance controller selectively randomly discarding packets having a traffic class association based on the current token availability level being within a token availability region specifying random packet discard of packets of the traffic class.

(Emphasis added). In rejecting independent claim 9, the Office Action maintains the rejections under the combination of *Carter*, *Patel*, *Gracon*, and *Lee*. The Office Action alleges that *Lee* discloses “c. a plurality of packet discard probability registers, each packet discard probability register specifying a probability with which packets of a specific traffic class are to be dropped when a current token availability level is within a token availability region.” In doing so, the Office Action continues to refer to the

registers 211 depicted in FIG. 2 and to the “drop probability” disclosed by *Lee* in FIG. 37 and to col. 56, lines 22-55, among other text passages. *Lee* teaches of drop probability used in determining whether to discard a packet and further teaches that the information element is discarded based on a drop probability and that the drop probability is calculated according to the equation: drop probability=((average ‘information segment storage unit’ occupancy-minimum number of occupied ‘information segment storage unit’ rows)/G)*(I).” (Col. 56, lines 43-47). *Lee*, however, fails to disclose or suggest the registers (211) specifying a probability with which packets of a specific traffic class are to be dropped when a current token availability level is within a token availability region. FIG. 2 shows the processing and context switching occurring in a prior art RISC processor performing networking functions. *Lee* teaches that processes (205) and (207) depicted in FIG. 2 use a common set of registers (211) to store information specific to that process. Nowhere does *Lee* appear to teach that the registers (211) are related to the drop probability described later in the disclosure. Furthermore, the *Carter*, *Patel*, and *Gracon* references fail to address this deficiency. Claim 9 is thus believed to be patentable.

2. Dependent Claims 10 and 12-15

Dependent claims 10 and 12-15 are allowable for at least the reason that these claims depend from an allowable independent claim. See, e.g., *In re Fine*, 837 F. 2d 1071 (Fed. Cir. 1988).

3. Independent Claim 24

The non-final Office Action mailed April 14, 2009 rejected claims 24-27 under 35 U.S.C. §103(a) as allegedly being unpatentable over *Carter* in view of *Patel*, and

Gracon in further view of *Lee*. For at least the reasons set forth below, Appellants submit that these rejections should be overturned.

Appellants' claim 24 provides as follows:

24. A method, of effecting ingress rate control comprising the step of:

selectively randomly discarding packets of a particular traffic class when a current token availability level of a leaky bucket tracking packets is between two token availability threshold levels of a plurality of token availability threshold levels, **wherein the token availability threshold levels correspond to predetermined ingress rate control responses** to bandwidth utilization with respect to packet traffic classes.

(Emphasis added). Claim 24 recites "wherein the token availability threshold levels correspond to predetermined ingress rate control responses." The Office Action equates these features with "*determining whether 'average usage of a class to which a flow belongs' is equal to, less than, or greater than a minimum/maximum threshold. [Figs. 1-2]*" as taught by *Lee*. (Office Action mailed April 14, 2009, page 20). The Office Action further refers to the accepting or discarding of incoming information elements depicted in FIGS. 34-41. Claim 24 explicitly recites "token availability threshold levels" in addition to "predetermined ingress rate control responses." Appellants respectfully submit that these elements are not taught by *Lee*.

The token availability threshold levels are used in the context of leaky bucket tracking packet transmissions ("when a current token availability level of a leaky bucket tracking packets is between two token availability threshold levels of a plurality of token availability threshold levels"). Even assuming, for the sake of argument, that the minimum and maximum thresholds cited in the Office Action correspond with the token availability threshold levels in claim 24, for example, *Lee* fails to disclose or suggest token availability threshold levels corresponding to predetermined ingress rate control

responses to bandwidth utilization. In FIG. 34, the compare unit 966 compares the average “information segment storage unit” occupancy of a particular class using the average occupancy counter for that class with the maximum number of occupied “information segment storage unit” rows and the minimum number of occupied “information segment storage unit” rows for that class. Based on this, the multiplexer 964 selects as its output a particular one of the inputs (e.g., “always discard”). Appellants submit that is not equivalent to the predetermined ingress rate control responses defined in claim 24.

4. Dependent Claims 25-27

Appellants submit that dependent claims 25-27 are allowable for at least the reason that these claims depend from an allowable independent claim. *See, e.g., In re Fine*, 837 F. 2d 1071 (Fed. Cir. 1988).

CONCLUSION

For at least the reasons discussed above, Appellants respectfully requests that the Examiner's rejection of claims 1-27 be overturned by the Board. In addition to the claims listed in Section VIII (CLAIMS – APPENDIX), Section IX (EVIDENCE – APPENDIX) included herein indicates that there is no additional evidence relied upon by this brief. Section X (RELATED PROCEEDINGS – APPENDIX) included herein indicates that there are no related proceedings.

A credit card authorization is provided to cover the fee with the Appeal Brief. No additional fee is believed to be due in connection with this appeal. If, however, any additional fee is deemed to be payable, you are hereby authorized to charge any such fee to deposit account 20-0778.

Respectfully submitted,

/ Jeffrey Hsu /

Jeffrey C. Hsu
Reg. No. 63,063

**THOMAS, KAYDEN, HORSTEMEYER
& RISLEY, L.L.P.**
600 Galleria Parkway S.E.
Suite 1500
Atlanta, Georgia 30339
(770) 933-9500

VIII. CLAIMS - APPENDIX

1. An egress rate controller monitoring content traffic transmitted from an edge network node of a packet-switched communications network node comprising:

a. a leaky bucket having an initial maximum number of tokens which decreases as packets are received in an associated output buffer at a reception token rate for transmission, wherein a size of the leaky bucket is less than or equal to a size of the associated output buffer;

b. a plurality of token availability threshold level registers specifying a corresponding plurality of token amounts defining token availability regions; and

c. a packet transmission suppression controller selectively suppressing transmission of a packet having a traffic class association based on a current token availability level being within a token availability region specifying transmission suppression of packets of the traffic class.

2. The egress rate controller claimed in claim 1, further comprising a classifier classifying received packets in accordance with a plurality of traffic classes.

3. The egress rate controller claimed in claim 1, further comprising a scheduler delaying packet transmission scheduling in accordance with a packet transmission suppression signal provided by the packet transmission suppression controller.

4. The egress rate controller claimed in claim 1, further comprising a bucket size register holding a value representative of a maximum number of tokens allocated to the leaky bucket.

5. The egress rate controller claimed in claim 4, further comprising an output buffer, the size of the leaky bucket, in tokens, being at most equal to the size of output buffer, employing an output buffer larger than the leaky bucket enabling suppression of packet transmission without discarding packets.

6. The egress rate controller claimed in claim 1, wherein the egress rate controller is associated with an output port of the edge network node.

7. An communication network node comprising at least one egress rate controller claimed in claim 1.

8. An communication network node comprising at least one egress rate controller claimed in claim 1 associated with at least one output port thereof.

9. An ingress rate controller monitoring content traffic received at an edge network node of a packet-switched communications network node comprising:
- a. a leaky bucket having an initial maximum number of tokens which decreases as packets received at a reception token rate are accepted;
 - b. a plurality of token availability threshold level registers specifying a corresponding plurality of token amounts defining token availability regions;
 - c. a plurality of packet discard probability registers, each packet discard probability register specifying a probability with which packets of a specific traffic class are to be dropped when a current token availability level is within a token availability region, and
 - d. a packet acceptance controller selectively randomly discarding packets having a traffic class association based on the current token availability level being within a token availability region specifying random packet discard of packets of the traffic class.
10. The ingress rate controller claimed in claim 9, further comprising a classifier classifying received packets in accordance with a plurality of traffic classes.
11. The ingress rate controller claimed in claim 9, further comprising a bucket size register holding a value representative of a maximum number of tokens allocated to the leaky bucket.
12. The ingress rate controller claimed in claim 9, further comprising an input buffer, the size of the leaky bucket, in tokens, being at most equal to the size of input

buffer, employing an input buffer larger than the leaky bucket providing a slack in the number of packets available for transmission to mask the effects of the ingress rate control effected.

13. The ingress rate controller claimed in claim 9, wherein the ingress rate controller is associated with an input port of the edge network node.

14. An communication network node comprising at least one ingress rate controller claimed in claim 9.

15. An communication network node comprising at least one ingress rate controller claimed in claim 9 associated with at least one input port thereof.

16. A method of effecting egress rate control comprising the step of:
selectively suppressing packet transmission for a packet of a particular traffic class when a current token availability level of a leaky bucket tracking packet transmissions is between two token availability threshold levels of a plurality of token availability threshold levels, wherein the token availability threshold levels correspond to predetermined egress rate control responses to bandwidth utilization with respect to packet traffic classes.

17. The method of effecting egress rate control as claimed in claim 16, wherein selectively suppressing packet transmission, the method further comprises a step of:

selectively suppressing packet transmission scheduling.

18. The method of effecting egress rate control as claimed in claim 17, further comprising a step of: rescheduling the packet for transmission.

19. The method of effecting egress rate control as claimed in claim 16, further comprising a prior step of: classifying packets in accordance with a plurality of traffic classes.

20. The method of effecting egress rate control as claimed in claim 16, further comprising a step of:

a. determining whether a plurality of tokens corresponding to a size of the packet are available in the leaky bucket; and

b. selectively suppressing packet transmission if there are insufficiently many tokens available in the leaky bucket.

21. The method of effecting egress rate control as claimed in claim 20, wherein selectively suppressing packet transmission, the method further comprises a step of: selectively suppressing packet transmission scheduling.

22. The method of effecting egress rate control as claimed in claim 21, further comprising a step of: storing the packet in an output buffer.

23. The method of effecting egress rate control as claimed in claim 21, further comprising a step of: rescheduling the packet for transmission.

24. A method, of effecting ingress rate control comprising the step of:
selectively randomly discarding packets of a particular traffic class when a current token availability level of a leaky bucket tracking packets is between two token availability threshold levels of a plurality of token availability threshold levels, wherein the token availability threshold levels correspond to predetermined ingress rate control responses to bandwidth utilization with respect to packet traffic classes.

25. The method of effecting ingress rate control as claimed in claim 24, wherein randomly discarding packets the method further comprises a step of: randomly discarding packets with a corresponding discard probability.

26. The method of effecting ingress rate control as claimed in claim 24, further comprising a prior step of: classifying packets in accordance with a plurality of traffic classes.

27. The method of effecting ingress rate control as claimed in claim 24, further comprising a step of:

- a. determining whether a plurality of tokens corresponding to a size of the packet are available in the leaky bucket; and
- b. selectively discarding the packet if there are insufficiently many tokens

available in the leaky bucket.

IX. EVIDENCE - APPENDIX

None.

X. RELATED PROCEEDINGS - APPENDIX

None.